

Exhibit 4

U.S. Patent No. 7,784,058 vs. HPE


VirtaMove asserts that Defendant Hewlett Packard Enterprise Company (“Defendant” or “HPE”) infringes the following claims (collectively, “Asserted Claims”): U.S. Patent No. 7,784,058 (“the ’058 patent”), claims 1–4 and 18.

Accused Instrumentalities: HPE products and services using user mode critical system elements as shared libraries, including without limitation HPE’s Ezmeral Runtime Enterprise (including without limitation both Ezmeral Runtime Enterprise and Ezmeral Runtime Enterprise Essentials, in each case including when marketed, sold, and/or licensed as part of or associated with HPE’s GreenLake branding, e.g. “HPE GreenLake for containers” which “is built on HPE Ezmeral Container Platform”), and all versions and variations thereof since the issuance of the asserted patent.


Each Accused Instrumentality infringes the claims in substantially the same way, and the evidence shown in this chart is similarly applicable to each Accused Instrumentality. Each claim limitation is literally infringed by each Accused Instrumentality. However, to the extent any claim limitation is not met literally, it is nonetheless met under the doctrine of equivalents because the differences between the claim limitation and each Accused Instrumentality would be insubstantial, and each Accused Instrumentality performs substantially the same function, in substantially the same way, to achieve the same result as the claimed invention. Notably, Defendant has not yet articulated which, if any, particular claim limitations it believes are not met by the Accused Instrumentalities.

Claim 1

Claim 1	Accused Instrumentalities
<p>[1pre] 1. A computing system for executing a plurality of software applications comprising:</p>	<p>To the extent the preamble is limiting, each Accused Instrumentality comprises or constitutes a computing system for executing a plurality of software applications as claimed.</p> <p><i>See</i> claim limitations below.</p> <p><i>See also, e.g.:</i></p> <p>HPE Ezmeral Runtime Enterprise is a unified platform built on open-source Kubernetes and designed for both cloud-native applications and non-cloud-native applications running on any infrastructure; whether on-premises, in multiple public clouds, in a hybrid model, or at the edge.</p> <p>https://support.hpe.com/hpsc/public/docDisplay?docId=a00ecp54hen_us&page=home/about-hpe-ezmeral-container-pl/Welcome.html</p>

Claim 1	Accused Instrumentalities
	<p>Containers provide the core runtime abstraction for the user applications. These containers provide isolation between user applications and the rest of the infrastructure. The containers are based on Docker.</p> <p>https://support.hpe.com/hpsc/public/docDisplay?docId=a00097165en_us&docLocale=en_US&page=GUID-6B6676DB-AF5F-4555-B6AB-D2C11A89F320.html</p> <p>Two Linux containers on a single system</p>  <pre> graph TD subgraph Container1 [Container #1 (OScontainer)] A1[Applications] B1[Bins/libs] C1[Init system] end subgraph Container2 [Container #2 (Appcontainer)] A2[Applications] end D[OS kernel] E[HW] </pre> <p>https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf</p>
[1a] a) a processor;	<p>Each Accused Instrumentality comprises a processor.</p> <p>For example, each node/host contains at least one CPU.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<p>Each license allows the customer to deploy the HPE Ezmeral Container Platform on one Core and 2 terabytes of Storage Capacity. The customer must purchase more licenses if they exceed the allowable amount of Cores or Storage Capacity. As used in this Agreement, Core means a part of a CPU that executes a single stream of compiled instruction code. Each physical processor contains smaller processing units called physical CPU cores. Some processors have two cores, some https://support.hpe.com/hpsc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_End_User_Software_Agreement.html</p>
<p>[1b] b) an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor; and,</p>	<p>Each Accused Instrumentality comprises an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor. For example, the OSCSEs include kernel-mode functions similar to the functionalities provided by user-space libraries such as glibc. These are implemented in kernel-space to handle tasks such as (without limitation) memory management (kmalloc(), kfree(), etc,) at kernel level.</p> <p><i>See, e.g.:</i></p> <p>Containers provide the core runtime abstraction for the user applications. These containers provide isolation between user applications and the rest of the infrastructure. The containers are based on Docker.</p> <p>https://support.hpe.com/hpsc/public/docDisplay?docId=a00097165en_us&docLocale=en_US&page=GUID-6B6676DB-AF5F-4555-B6AB-D2C11A89F320.html</p> <p>Each license allows the customer to deploy the HPE Ezmeral Container Platform on one Core and 2 terabytes of Storage Capacity. The customer must purchase more licenses if they exceed the allowable amount of Cores or Storage Capacity. As used in this Agreement, Core means a part of a CPU that executes a single stream of compiled instruction code. Each physical processor contains smaller processing units called physical CPU cores. Some processors have two cores, some https://support.hpe.com/hpsc/public/docDisplay?docId=a00ecp54hen_us&docLocale=en_US&page=home/about-hpe-ezmeral-container-pl/GEN_End_User_Software_Agreement.html</p>


Claim 1	Accused Instrumentalities				
	<p>HPE Ezmeral Runtime Enterprise supports the following operating systems:</p> <table border="0" data-bbox="713 295 1860 474"> <tr> <td data-bbox="713 295 946 474">HPE Ezmeral Runtime Enterprise Version</td> <td data-bbox="1030 295 1262 474">CentOS Support</td> <td data-bbox="1262 295 1579 474">RHEL Support</td> <td data-bbox="1579 295 1860 474">SUSE Support</td> </tr> </table> <p>https://support.hpe.com/hpsc/public/docDisplay?docId=a00ecp54hen_us&page=home/about-hpe-ezmeral-container-pl/GEN_OS_Support.html</p> <p>Two Linux containers on a single system</p>  <pre> graph TD subgraph Container1 [Container #1 (OScontainer)] direction TB A1[Applications] --- B1[Bins/libs] B1 --- C1[Init system] C1 --- D1[OS kernel] D1 --- E1[HW] end subgraph Container2 [Container #2 (Appcontainer)] direction TB A2[Applications] end </pre> <p>https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA-6-2761ENW.pdf</p>	HPE Ezmeral Runtime Enterprise Version	CentOS Support	RHEL Support	SUSE Support
HPE Ezmeral Runtime Enterprise Version	CentOS Support	RHEL Support	SUSE Support		

Claim 1	Accused Instrumentalities
	<p>Kernel mode</p> <p>Kernel mode refers to the processor mode that enables software to have full and unrestricted access to the system and its resources. The OS kernel and kernel drivers, such as the file system driver, are loaded into protected memory space and operate in this highly privileged kernel mode.</p> <p>https://www.techtarget.com/searchdatacenter/definition/kernel</p> <p>The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p>https://en.wikipedia.org/wiki/Glibc</p>
[1c] c) a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode and	<p>Each Accused Instrumentality comprises a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode. For example, the shared library with SLCSEs include the runtime environment, system tools, and dependencies, such as the glibc library and other libraries that replicate OSCSEs, included in the container image (including without limitation in a base image that is included within the container image).</p> <p><i>See, e.g.:</i></p> <p>Containers provide the core runtime abstraction for the user applications. These containers provide isolation between user applications and the rest of the infrastructure. The containers are based on Docker.</p> <p>https://support.hpe.com/hpsc/public/docDisplay?docId=a00097165en_us&docLocale=en_US&page=GUID-6B6676DB-AF5F-4555-B6AB-D2C11A89F320.html</p> <p>The container starts with a base image, and the microservice is packaged into a container image and then deployed through the container platform. The container platform is based on</p> <p>https://www.hpe.com/us/en/what-is/container-platform.html</p>

Claim 1	Accused Instrumentalities
	<p>Container images</p> <p>A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <hr/> <p>https://kubernetes.io/docs/concepts/containers/</p> <p>The idea of containerization is to isolate and package the application with all the dependencies in a container.</p> <p>https://community.hpe.com/t5/hpe-blog-uk-ireland-middle-east/containerization-the-next-generation-of-virtualization/ba-p/7154442</p> <p>Container image files are complete, static and executable versions of an application or service and differ from one technology to another. Docker images are made up of multiple layers, which start with a base image that includes all of the dependencies needed to execute code in a container. Each image has a readable/writable layer on top of static unchanging layers. Because each container has its own specific container layer that customizes that specific container, underlying image layers can be saved and reused in multiple containers. An Open Container Initiative (OCI)</p> <p>https://www.techtarget.com/searchitoperations/definition/container-containerization-or-container-based-virtualization</p>

Claim 1	Accused Instrumentalities
	<h2>About storage drivers</h2> <p>To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2>Storage drivers versus Docker volumes</h2> <p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the volumes section to learn how to use volumes to persist data and improve performance.</p> <p>https://docs.docker.com/storage/storagedriver/</p>


Claim 1	Accused Instrumentalities
	<h2 data-bbox="692 208 1136 262">Images and layers</h2> <p data-bbox="692 300 1875 376">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="713 442 1501 763"># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="692 829 1959 1139">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="671 1160 1284 1192">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>Container (based on <code>ubuntu:15.04</code> image)</p> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h1>Volumes</h1> <p>Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p>https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2>Container environment</h2> <p>The Kubernetes Container environment provides several important resources to Containers:</p> <ul style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p>https://kubernetes.io/docs/concepts/containers/container-environment/</p>


Claim 1	Accused Instrumentalities
	<h2>Images</h2> <p>A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p>You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p>https://kubernetes.io/docs/concepts/containers/images/</p>
	<h2>Volumes</h2>
	<p>On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p>
	<p>https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<p>Open Container Initiative</p> <hr/> <p>Image Format Specification</p> <hr/> <p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p> <p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<p>Overview</p> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p>  <pre data-bbox="696 718 1034 816"> public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } } </pre> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<h2>OCI Image Configuration</h2> <p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p>Layer</p> <ul style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p>Image JSON</p> <ul style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="726 213 1410 251">Two Linux containers on a single system</p>  <pre> graph TD subgraph Container1 [Container #1 (OScontainer)] A1[Applications] B1[Bins/libs] C1[Init system] end subgraph Container2 [Container #2 (Appcontainer)] A2[Applications] end OS[OS kernel] HW[HW] Container1 --- OS Container2 --- OS Container1 --- HW Container2 --- HW </pre> <p data-bbox="677 931 1981 1002">https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf</p> <p data-bbox="677 1039 1966 1204">The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p data-bbox="677 1220 1136 1253">https://en.wikipedia.org/wiki/Glibc</p>
[1d] i) wherein some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when one of the	In each Accused Instrumentality, some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications.


Claim 1	Accused Instrumentalities
<p>SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications,</p>	<p>For example, a base image serves as a self-contained unit that encompasses all the necessary components for an application to run, including the application code, runtime environment, system tools, and dependencies (i.e., SLCSEs). The images are based on existing Linux distributions, such as Debian and Ubuntu, including essential system elements (i.e., functional replicas of OSCSEs). Each container image is based on a specific base image, which contains the application code, and dependencies, including system libraries or shared library critical system elements (SLCSEs). The base image forms a part of the container image according to the “layer” model described in the documentation below. When the container runs the image, it creates a runtime instance of that container image. In turn, when one or more applications executes within the container runtime environment, it dynamically links to the SLCSEs stored in the runtime environment, which thereby become a part of the application(s).</p> <p><i>See, e.g.:</i></p> <p>Hewlett Packard Enterprise provides publicly available base OS images for use in containerized clusters. These images extend the base OS images available from Docker hub by adding several packages that permit HPE Ezmeral Runtime Enterprise to manage container orchestration seamlessly and to improve the security of the container.</p> <p>https://docs.ezmeral.hpe.com/runtime-enterprise/55/app-workbench-5-1/custom-base-images/AWB51_About_Custom_Base_Images.html</p> <p>The idea of containerization is to isolate and package the application with all the dependencies in a container.</p> <p>https://community.hpe.com/t5/hpe-blog-uk-ireland-middle-east/containerization-the-next-generation-of-virtualization/ba-p/7154442</p> <h2>Container images</h2> <p>A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p>https://kubernetes.io/docs/concepts/containers/</p>

Claim 1	Accused Instrumentalities																																																						
	<div data-bbox="692 208 1945 425">  <p>ubuntu  Updated 15 days ago Ubuntu is a Debian-based Linux operating system based on free software. Linux IBM Z 386 riscv64 x86-64 ARM ARM 64 PowerPC 64 LE</p> </div>																																																						
	<div data-bbox="692 470 1945 687">  <p>debian  Updated 35 minutes ago Debian is a Linux distribution that's composed entirely of free and open-source software. Linux riscv64 x86-64 ARM ARM 64 386 mips64le PowerPC 64 LE IBM Z</p> </div>																																																						
<p>https://hub.docker.com/search?image_filter=official&type=image&q=</p>																																																							
	<table border="1"> <thead> <tr> <th>Platform</th> <th>x86_64 / amd64</th> <th>arm64 / aarch64</th> <th>arm (32-bit)</th> <th>ppc64le</th> <th>s390x</th> </tr> </thead> <tbody> <tr> <td>CentOS</td> <td>✓</td> <td>✓</td> <td></td> <td>✓</td> <td></td> </tr> <tr> <td>Debian</td> <td>✓</td> <td>✓</td> <td>✓</td> <td>✓</td> <td></td> </tr> <tr> <td>Fedora</td> <td>✓</td> <td>✓</td> <td></td> <td>✓</td> <td></td> </tr> <tr> <td>Raspberry Pi OS (32-bit)</td> <td></td> <td></td> <td>✓</td> <td></td> <td></td> </tr> <tr> <td>RHEL (s390x)</td> <td></td> <td></td> <td></td> <td>✓</td> <td></td> </tr> <tr> <td>SLES</td> <td></td> <td></td> <td></td> <td>✓</td> <td></td> </tr> <tr> <td>Ubuntu</td> <td>✓</td> <td>✓</td> <td>✓</td> <td>✓</td> <td>✓</td> </tr> <tr> <td>Binaries</td> <td>✓</td> <td>✓</td> <td>✓</td> <td></td> <td></td> </tr> </tbody> </table>	Platform	x86_64 / amd64	arm64 / aarch64	arm (32-bit)	ppc64le	s390x	CentOS	✓	✓		✓		Debian	✓	✓	✓	✓		Fedora	✓	✓		✓		Raspberry Pi OS (32-bit)			✓			RHEL (s390x)				✓		SLES				✓		Ubuntu	✓	✓	✓	✓	✓	Binaries	✓	✓	✓		
Platform	x86_64 / amd64	arm64 / aarch64	arm (32-bit)	ppc64le	s390x																																																		
CentOS	✓	✓		✓																																																			
Debian	✓	✓	✓	✓																																																			
Fedora	✓	✓		✓																																																			
Raspberry Pi OS (32-bit)			✓																																																				
RHEL (s390x)				✓																																																			
SLES				✓																																																			
Ubuntu	✓	✓	✓	✓	✓																																																		
Binaries	✓	✓	✓																																																				

Claim 1	Accused Instrumentalities
	<p><u>https://docs.docker.com/engine/install/</u></p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p><u>https://www.techtarget.com/searchitoperations/definition/Docker-image</u></p>

Claim 1	Accused Instrumentalities
	<h2>About storage drivers</h2> <p>To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2>Storage drivers versus Docker volumes</h2> <p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the volumes section to learn how to use volumes to persist data and improve performance.</p> <p>https://docs.docker.com/storage/storagedriver/</p>


Claim 1	Accused Instrumentalities
	<h2 data-bbox="692 208 1136 264">Images and layers</h2> <p data-bbox="692 300 1875 376">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="713 442 1501 763"># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="692 829 1959 1139">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="671 1160 1284 1192">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p>  <p>Container (based on <code>ubuntu:15.04</code> image)</p> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h1>Volumes</h1> <p>Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p>https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2>Container environment</h2> <p>The Kubernetes Container environment provides several important resources to Containers:</p> <ul style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p>https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2>Images</h2> <p>A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p>You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p>https://kubernetes.io/docs/concepts/containers/images/</p>
	<h2>Volumes</h2>
	<p>On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p>
	<p>https://kubernetes.io/docs/concepts/storage/volumes/</p>


Claim 1	Accused Instrumentalities
	<p>Open Container Initiative</p> <hr/> <p>Image Format Specification</p> <hr/> <p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p> <p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>


Claim 1	Accused Instrumentalities
	<h2>Overview</h2> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p>  <pre data-bbox="713 731 1036 822"> public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } } </pre> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>


Claim 1	Accused Instrumentalities
	<h2>OCI Image Configuration</h2> <p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p>Layer</p> <ul style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p>Image JSON</p> <ul style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p> <p>Containers only have access to resources that are defined in the image, https://www.hpe.com/us/en/what-is/docker.html</p>

Claim 1	Accused Instrumentalities
<p>[1e] ii) wherein an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function, and</p>	<p>DESCRIPTION top</p> <p>The programs ld.so and ld-linux.so* find and load the shared objects (shared libraries) needed by a program, prepare the program to run, and then run it.</p> <p>https://man7.org/linux/man-pages/man8/ld.so.8.html</p> <p>In each Accused Instrumentality, an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function.</p> <p>When a Docker or Kubernetes image is used to create a container, it creates a separate and isolated instance of a runtime environment which is independent of other containers running on the same host. Each container has its own instance of base images and its own data. The containers run in isolation, ensuring that the SLCSEs stored in the shared library are accessible to the software applications running in their respective containers. The image includes essential system files, libraries, and dependencies required to run the software application within the container. The containers can share common dependencies and components using layered images. This means that multiple containers utilize the same base image to create an instance. When an instance of SLCSE is provided from the base image (i.e., from the shared library) to an individual container including application software, it operates in isolation and runs its own instance of the software application without sharing resources or critical system elements with other containers. This ensures that each container has its own isolated context. Docker or Kubernetes containers can share common dependencies and components using layered images. This means that multiple containers can utilize the same base image. Therefore, each container, containing the application software running under the operating system, utilizes a unique instance of the corresponding critical system element to execute the respective application software for performing a same function.</p> <p><i>See, e.g.:</i></p>


Claim 1	Accused Instrumentalities
	<p>Containers provide the core runtime abstraction for the user applications. These containers provide isolation between user applications and the rest of the infrastructure. The containers are based on Docker.</p> <p>https://support.hpe.com/hpsc/public/docDisplay?docId=a00097165en_us&docLocale=en_US&page=GUID-6B6676DB-AF5F-4555-B6AB-D2C11A89F320.html</p> <p>Two Linux containers on a single system</p>  <pre> graph TD subgraph Container1 [Container #1 (OS container)] A1[Applications] B1[Bins/libs] C1[Init system] end subgraph Container2 [Container #2 (App container)] A2[Applications] end OS[OS kernel] HW[HW] Container1 --- OS Container1 --- HW Container2 --- OS Container2 --- HW </pre> <p>https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf</p>

Claim 1	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image</p>
<p>[1f] iii) wherein a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p>	<p>In each Accused Instrumentality, a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p> <p>For example, in Docker or Kubernetes containers, each container operates independently, and a base image includes essential system files, libraries, and dependencies (i.e., SLCSEs) required to run the software application within the container. Based on information and belief, each element, such as system files, libraries, and dependencies (i.e., SLCSE) is associated with an execution of a predetermined function related to the application. When an image is used to create a container in the Accused Instrumentality, an instance of the SLCSE is provided to a software application. Therefore,</p>

Claim 1	Accused Instrumentalities
	<p>different instances of the SLCSE are provided to different applications for performing a same function, simultaneously.</p> <p><i>See, e.g.:</i></p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image, Last accessed on June 14, 2023</p> <p>A container is a runnable instance of an image. You can create, start, stop, move, or delete a container using the Docker API or CLI. You can connect a container to one or more networks, attach storage to it, or even create a new image based on its current state.</p> <p>https://docs.docker.com/get-started/overview/</p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 2

Claim 2	Accused Instrumentalities
<p>2. A computing system as defined in claim 1, wherein in operation, multiple instances of an SLCSE stored in the shared library run simultaneously within the operating system.</p>	<p>Each Accused Instrumentality comprises or constitutes a computing system as defined in claim 1, wherein in operation, multiple instances of an SLCSE stored in the shared library run simultaneously within the operating system.</p> <p>For example, an individual host/node runs multiple containers and/or pods simultaneously, each of which has an instance of an SLCSE. When the multiple containers and/or pods run simultaneously, the multiple instances of the SLCSE stored in the shared library run simultaneously.</p> <p><i>See, e.g.:</i></p>



Claim 2	Accused Instrumentalities
	<h1>Containers</h1> <p>Each container that you run is repeatable; the standardization from having dependencies included means that you get the same behavior wherever you run it.</p> <p>Containers decouple applications from the underlying host infrastructure. This makes deployment easier in different cloud or OS environments.</p> <p>Each <u>node</u> in a Kubernetes cluster runs the containers that form the <u>Pods</u> assigned to that node. Containers in a Pod are co-located and co-scheduled to run on the same node.</p> <p>https://kubernetes.io/docs/concepts/containers/</p>


Claim 2	Accused Instrumentalities
	<h1>Kubernetes Scheduler</h1> <p>In Kubernetes, <i>scheduling</i> refers to making sure that Pods are matched to Nodes so that Kubelet can run them.</p> <h2>Scheduling overview</h2> <p>A scheduler watches for newly created Pods that have no Node assigned. For every Pod that the scheduler discovers, the scheduler becomes responsible for finding the best Node for that Pod to run on. The scheduler reaches this placement decision taking into account the scheduling principles described below.</p> <p>If you want to understand why Pods are placed onto a particular Node, or if you're planning to implement a custom scheduler yourself, this page will help you learn about scheduling.</p> <p>https://kubernetes.io/docs/concepts/scheduling-eviction/kube-scheduler/</p>


Claim 2	Accused Instrumentalities
	<h1>Running containers</h1> <p>Docker runs processes in isolated containers. A container is a process which runs on a host. The host may be local or remote. When you execute <code>docker run</code>, the container process that runs is isolated in that it has its own file system, its own networking, and its own isolated process tree separate from the host.</p> <p>https://docs.docker.com/engine/reference/run/</p>

Claim 3

Claim 3	Accused Instrumentalities
<p>3. A computing system according to claim 1 wherein OSCSEs corresponding to and capable of performing the same function as SLCSEs remain in the operating system kernel.</p>	<p>Each Accused Instrumentality comprises or constitutes a computing system according to claim 1 wherein OSCSEs corresponding to and capable of performing the same function as SLCSEs remain in the operating system kernel.</p> <p>For example, both Docker and Kubernetes systems preserve the host kernel substantially unchanged; therefore the OSCSEs corresponding to the SLCSEs remain in the operating system kernel.</p> <p><i>See, e.g.:</i></p> <p>Containers provide the core runtime abstraction for the user applications. These containers provide isolation between user applications and the rest of the infrastructure. The containers are based on Docker.</p> <p>https://support.hpe.com/hpsc/public/docDisplay?docId=a00097165en_us&docLocale=en_US&page=GUID-6B6676DB-AF5F-4555-B6AB-D2C11A89F320.html</p> <p>The container starts with a base image, and the microservice is packaged into a container image and then deployed through the container platform. The container platform is based on https://www.hpe.com/us/en/what-is/container-platform.html</p>

Claim 3	Accused Instrumentalities
	<p>Container images</p> <p>A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <hr/> <p>https://kubernetes.io/docs/concepts/containers/</p> <p>The idea of containerization is to isolate and package the application with all the dependencies in a container.</p> <p>https://community.hpe.com/t5/hpe-blog-uk-ireland-middle-east/containerization-the-next-generation-of-virtualization/ba-p/7154442</p> <p>Container image files are complete, static and executable versions of an application or service and differ from one technology to another. Docker images are made up of multiple layers, which start with a base image that includes all of the dependencies needed to execute code in a container. Each image has a readable/writable layer on top of static unchanging layers. Because each container has its own specific container layer that customizes that specific container, underlying image layers can be saved and reused in multiple containers. An Open Container Initiative (OCI)</p> <p>https://www.techtarget.com/searchitoperations/definition/container-containerization-or-container-based-virtualization</p>

Claim 3	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 3	Accused Instrumentalities
<p style="text-align: center;">Two Linux containers on a single system</p>  <pre> graph TD subgraph Container1 [Container #1 (OScontainer)] A1[Applications] B1[Bins/libs] C1[Init system] end subgraph Container2 [Container #2 (Appcontainer)] A2[Applications] end D[OS kernel] E[HW] </pre> <p>https://h50146.www5.hpe.com/products/software/oe/linux/mainstream/support/whitepaper/pdfs/4AA6-2761ENW.pdf</p>	

Claim 4


Claim 4	Accused Instrumentalities
<p>4. A computing system according to claim 1 wherein the one or more SLCSEs provided to one of the plurality of software applications having exclusive use thereof, use system calls to access services in the operating system kernel.</p>	<p>Each Accused Instrumentality comprises or constitutes a computing system according to claim 1 wherein the one or more SLCSEs provided to one of the plurality of software applications having exclusive use thereof, use system calls to access services in the operating system kernel.</p> <p>For example, the SLCSEs in a container use system calls to access services in the operating system kernel. For example, the glibc library (or other similar library) in the container uses system calls to</p>

Claim 4	Accused Instrumentalities
	<p>interface with the host Linux kernel. In general, system calls can be observed using a tool such as strace.</p> <p><i>See, e.g.:</i></p> <p>The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p>https://en.wikipedia.org/wiki/Glibc</p>

Claim 4	Accused Instrumentalities
	<p>We can now get the process id directly from the cgroup. It will be located in the cgroup.procs file.</p> <pre>### Terminal 2 - Worker Node ### # Get the process id \$ cat cri-containerd-ceeeef06afe89c8223d33b11e8d9e0b207118ac4dac3af826687668ee1ee 16254 # Validate what is running under the process \$ ps aux grep 16254 azureus+ 16254 0.0 0.1 713972 10476 ? Ssl 15:04 0:00 ./faultyapp azureus+ 94806 0.0 0.0 7004 2168 pts/0 S+ 16:22 0:00 grep --color=a</pre> <p>Got it! With that, we can try to find out what is going on inside the app. Let's try to run strace to get some more insight.</p> <pre>### Terminal 2 - Worker Node ### \$ sudo strace -p 16254 -f ... # The app is trying to read a file port.txt [pid 16269] openat(AT_FDCWD, "port.txt", O_RDONLY O_CLOEXEC <unfinished ...> [pid 16254] epoll_pwait(5, <unfinished ...> # The file does not exist [pid 16269] <... openat resumed> = -1 ENOENT (No such file or directory) [pid 16254] <... epoll_pwait resumed>[], 128, 0, NULL, 0) = 0 [pid 16269] write(1, "Something went wrong...\n", 24 <unfinished ...></pre> <p>After filtering the output, we can see the application is trying to read a text file called port.txt, and a few lines later, there is a message stating ENOENT (No such file or directory). Let's create that file.</p> <p>https://www.berops.com/blog/a-different-method-to-debug-kubernetes-pods</p>

Claim 18

Claim 18	Accused Instrumentalities
<p>18. A computer system as defined in claim 2 wherein SLCSEs are not copies of OSCSEs.</p>	<p>Each Accused Instrumentality comprises or constitutes a computer system as defined in claim 2 wherein SLCSEs are not copies of OSCSEs.</p> <p>For example, in a typical case the SLCSEs come from a Linux distribution independent of the host operating system, and thus are not identical to the OSCSEs.</p> <p><i>See, e.g.:</i></p> <p>Hewlett Packard Enterprise provides publicly available base OS images for use in containerized clusters. These images extend the base OS images available from Docker hub by adding several packages that permit HPE Ezmeral Runtime Enterprise to manage container orchestration seamlessly and to improve the security of the container.</p> <p>https://docs.ezmeral.hpe.com/runtime-enterprise/55/app-workbench-5-1/custom-base-images/AWB51_About_Custom_Base_Images.html</p> <p>The idea of containerization is to isolate and package the application with all the dependencies in a container.</p> <p>https://community.hpe.com/t5/hpe-blog-uk-ireland-middle-east/containerization-the-next-generation-of-virtualization/ba-p/7154442</p> <h2 data-bbox="713 915 1115 964">Container images</h2> <p>A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p>https://kubernetes.io/docs/concepts/containers/</p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p>

Claim 18	Accused Instrumentalities										
	<p data-bbox="671 200 1586 230">https://www.techtarget.com/searchitoperations/definition/Docker-image</p> <p data-bbox="686 282 1959 409">Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <table border="1" data-bbox="1136 665 1537 1029"> <tbody> <tr> <td>91e54dfb1179</td> <td>0 B</td> </tr> <tr> <td>d74508fb6632</td> <td>1.895 KB</td> </tr> <tr> <td>c22013c84729</td> <td>194.5 KB</td> </tr> <tr> <td>d3a1f33e8a5a</td> <td>188.1 MB</td> </tr> <tr> <td colspan="2">ubuntu:15.04</td> </tr> </tbody> </table> <p data-bbox="671 1062 1275 1095">https://docs.docker.com/storage/storagedriver/</p> <p data-bbox="671 1119 1320 1148">Containers only have access to resources that are defined in the image,</p> <p data-bbox="671 1160 1296 1191">https://www.hpe.com/us/en/what-is/docker.html</p>	91e54dfb1179	0 B	d74508fb6632	1.895 KB	c22013c84729	194.5 KB	d3a1f33e8a5a	188.1 MB	ubuntu:15.04	
91e54dfb1179	0 B										
d74508fb6632	1.895 KB										
c22013c84729	194.5 KB										
d3a1f33e8a5a	188.1 MB										
ubuntu:15.04											